

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Appellant(s):	Santosh S. Rao, Gopal Sharma, Poonam Dhavale		
Assignee:	Veritas Operating Corporation		
Title:	System And Method For Resolving Cluster Partitions In Out-Of-Band Storage Virtualization Environments		
Serial No.:	10/627,385	Filing Date:	July 25, 2003
Examiner:	Mohamed A. Wasel	Technology Center:	2100
Docket No.:	VRT0089US	Appeal No.	2008-0258

Austin, Texas
August 21, 2008

Board of Patent Appeals and Interference
US Patent and Trademark Office
PO Box 1450
Alexandria, Virginia 22313-1450

**SUPPLEMENTAL APPEAL BRIEF IN RESPONSE TO
ORDER REQUIRING APPELLANTS TO BRIEF AN ADDITIONAL
MATTER**

Dear Sirs:

This brief is being submitted in response to the Board of Patent Appeal and Interferences order requiring Appellants to brief an additional matter, to assist the Board in reaching a reasoned decision on the pending appeal. This brief includes a revised Summary of Claimed Subject Matter that is responsive to the order.

As the fee for the Appeal Brief has already been charged, no additional fees are believed due. Nevertheless, please charge deposit account 502306 for any additional sums which may be required to be paid as part of this appeal.

REAL PARTY IN INTEREST

The real party in interest on this appeal is Veritas Operating Corporation.

RELATED APPEALS AND INTERFERENCES

There are no appeals or interferences related to this application.

STATUS OF CLAIMS

Claims 1-31 stand rejected under 35 U.S.C. § 102(e) as being anticipated by Frank, et al., U.S. Patent No. 6,532,494 (Frank). Claims 1-31 are being appealed.

STATUS OF AMENDMENTS

No amendments were filed subsequent to the final rejection of January 27, 2006.

SUBSTITUTE SUMMARY OF CLAIMED SUBJECT MATTER

The invention is as set forth in the claims. To summarize the invention without intending to limit or otherwise affect the scope of the claims, the invention as set forth by independent claim 1 relates to a method. A coordinator virtual device (e.g., coordinator volume 400 of **Figure 4**) corresponding to at least a portion of a physical data storage device (e.g., SAN 160 of **Figure 2** or one or more of storage devices 170 – 190 thereof) is provided. See, for example, step 510 of **Figure 5** and paragraph 0051, lines 4-8. Please note paragraph 0031, lines 2-3 states, “In general, a volume is a virtual device or disk representing an addressable range of disk blocks ...” (Emphasis added.) Claim 1 requires detecting when a computer system cluster, including a plurality of nodes (e.g., cluster nodes 210 and 250), is partitioned. While the figures do not expressly show this step, the flow chart of **Figure 9** includes step 930, “cluster partition occurs.” See paragraph 0067, lines 1-3. Appellants assert the step of detecting when a computer system cluster is partitioned is supported in the detailed description. For example, paragraph 0043, lines 20-23 explains that “when node membership and messaging 310 informs other software components that communication with another cluster node has been lost, i.e., that the cluster has partitioned, fence driver 320 intercepts this message and begins taking action to protect shared resources.” Claim 1 requires an attempt is made to gain control of the coordinator virtual device. See, for example, step 935 or step 955 of **Figure 9** and paragraph 0067, lines 3-9. At least one of the plurality of nodes from the computer system cluster is removed when the attempting is unsuccessful. See, for example, ejection of step 945 of **Figure 9** and paragraph 0068, lines 4-5. As an

aside, paragraph 0047, lines 9-10 states that “ejected systems cannot write to the data disks and corrupt data.”

The invention as set forth by independent claim 15 relates to a system comprising a first data storage device (e.g., SAN 160 of **Figure 2** or one or more of storage devices 170 – 190 thereof), a virtual device configuration server (e.g., volume server 280 of **Figure 2**), and a plurality of virtual device configuration clients (e.g., cluster nodes 210 and 250 of **Figure 2**). Claim 15 requires the virtual device configuration server (e.g., volume server 280) to be coupled to the first storage device and includes a first memory and a first processor configured to provide a coordinator virtual device (e.g., coordinator volume 400 of **Figure 4**) corresponding to at least a portion of the first data storage device. The volume server 280 of **Figure 2** is not shown with a first memory and a first processor. Even though **Figure 2** does not show a first memory and a first processor, the detailed description supports the first memory and the first processor. For example, **Figure 5** shows creation of the coordinator volume at step 510, and paragraph 0075, lines 5-7 state, “Moreover, the methods described in **Figures 5-10** are typically implemented as one or more software programs for a computer system and are encoded in a computer readable medium as instructions executable on one or more processors.” (Emphasis added.) Claim 15 requires the plurality of virtual device configuration clients (e.g., cluster nodes 210 and 250) to be configured as a computer system cluster. See, for example, paragraph 0030, lines 2-3. Claim 15 requires at least one of the plurality of virtual device configuration clients to include a second memory and a second processor. Example support for the second memory is found in memory 300 of **Figure 3** and paragraph 0043 lines 1-5. Neither **Figure 2** nor **Figure 3** shows a second processor, but the detailed description supports the second processor of one of the virtual device configuration clients (e.g., cluster node 210 or 250). For example, paragraph 0077 provides:

Figure 11 illustrates a block diagram of a computer system 1100 for implementing the I/O strategy of the present invention. For example, computer system 1100 can be an embodiment of one of the previously described cluster nodes, client computer systems, host computer systems, or even network and storage appliances. Computer system 1100 includes a processor 1110 and a memory 1120 coupled together by communications bus 1105. Processor 1110 can be a single processor or a

number of individual processors working together. Memory 1120 is typically random access memory (RAM), or some other dynamic storage device, and is capable of storing instructions to be executed by the processor, e.g., application software 360 and fence driver 320. Memory 1120 is also used for storing temporary variables or other intermediate information during the execution of instructions by the processor 1110.

(Emphasis added.) Further, paragraph 0043 presumes software components of memory 300 of **Figure 3**, including fence driver 320, are executed on one or more processors. In Claim 15, one of the virtual device configuration clients (e.g., cluster node 210 or 250) can detect when the computer system cluster is partitioned. While the figures do not expressly show this step, the flow chart of **Figure 9** includes step 930, “cluster partition occurs.” See paragraph 0067, lines 1-3. Appellants assert the step of detecting when a computer system cluster... is partitioned is supported in the detailed description. For example, paragraph 0043, lines 20-23 explains that “when node membership and messaging 310 informs other software components that communication with another cluster node has been lost, i.e., that the cluster has partitioned, fence driver 320 intercepts this message and begins taking action to protect shared resources.” Claim 15 recites that an attempt can be made to gain control of the coordinator virtual device. See, for example, step 935 or step 955 of **Figure 9** and paragraph 0067, lines 3-9. At least one of the plurality of the virtual configuration clients (e.g., cluster node 210 or 250) is removed when the attempting is unsuccessful. See, for example, ejection of step 945 of **Figure 9** and paragraph 0068, lines 4-5. As an aside, paragraph 0047, lines 9-10 states that “ejected systems cannot write to the data disks and corrupt data.”

The invention as set forth by independent claim 28 relates to an apparatus. The apparatus includes a means (e.g., volume server 280 of **Figure 2**) for providing a coordinator virtual device (e.g., coordinator volume 400 of **Figure 4**). See, for example, paragraph 0051, lines 4-8. The coordinator virtual device corresponds to at least a portion of a physical data storage device (e.g., SAN 160 of **Figure 2** or one or more of storage devices 170 – 190 thereof). See, for example, paragraph 0051, lines 7-8. Further, paragraph 0031, lines 2-3 states, “ In general, a volume is a virtual device or disk representing an addressable range of disk blocks ...” (Emphasis added.) A means (e.g., fence driver 320 of **Figure 2**) for detecting when a computer system cluster, including a plurality of nodes (e.g., cluster nodes 210 and 250 of **Figure 2**), is partitioned is also

included. See, for example, paragraph 0043, lines 20-23. The system further includes a means (e.g., fence driver 320) for attempting to gain control of the coordinator virtual device. See, for example, paragraph 0067, lines 3-9. Finally, the system provides a means for removing at least one of the plurality of nodes from the computer system cluster when the attempting is unsuccessful. See, for example, paragraph 0068, lines 4-5.

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

- I. Claims 1-31 stand rejected under 35 U.S.C. § 102(e) as being anticipated by Frank, et al., U.S. Patent No. 6,532,494 (Frank).

ARGUMENT

Claims 1-31 stand rejected under 35 U.S.C. § 102(e) as being anticipated by Frank, et al., U.S. Patent No. 6,532,494 (Frank). Claims 1-31 are being appealed.

Frank fails to teach or suggest a method including:

providing a coordinator virtual device corresponding to at least a portion of a physical data storage device;

detecting when a computer system cluster, including a plurality of nodes, is partitioned;

attempting to gain control of the coordinator virtual device; and

removing at least one of the plurality of nodes from the computer system cluster when the attempting is unsuccessful,

as required by independent claim 1 and generally required by independent claim 28.

Regarding the claimed “providing a coordinator virtual device corresponding to at least a portion of a physical data storage device,” (and the related operation of attempting to gain control of the virtual coordinator device) the Examiner has identified no less than *four* different possible elements in Frank as teaching claimed coordinator virtual device.

These possible elements are:

- (1) a single virtual computer formed from cluster nodes and appearing to an end user as a single computing resource (Frank, column 1, lines 30-40 as cited in the Final Office Action of January 27, 2006, p. 2, no. 1(a));

(2) shareable storage 22 (Frank, column 3, lines 35-45 as cited in the Final Office Action of January 27, 2006, p. 2, no. 1(a));

(3) a cluster manager 32 (See, e.g., Frank, abstract and column 4, lines 31-42, as generally referenced in the Final Office Action of January 27, 2006, p. 12, ¶b); and

(4) a coordinator node (Frank, column 8, lines 43-61 as cited in the Advisor Action of April 24, 2006, p. 2, response ¶a).

As an initial matter, the appellants respectfully submit that such use of Frank is inappropriate as a basis for a §102 rejection. MPEP §2131 makes clear the requirements for anticipation:

“A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference.” *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987). . . . “The identical invention must be shown in as complete detail as is contained in the . . . claim.” *Richardson v. Suzuki Motor Co.*, 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989). The elements must be arranged as required by the claim, but this is not an *ipsissimis verbis* test, i.e., identity of terminology is not required. *In re Bond*, 910 F.2d 831, 15 USPQ2d 1566 (Fed. Cir. 1990). (Emphasis added)

Thus, in addition to showing every element, the reference must teach their arrangement as required by the claim. It is unclear how this requirement can be met by the Examiner when the Examiner has identified at least four distinct elements that might be the claimed coordinator virtual device. Nevertheless, each of these arguments are addressed below.

Frank’s single virtual computer formed from cluster nodes and appearing to an end user as a single computing resource is not a coordinator virtual device. The cited portion of Frank states:

As is known in the art, a computer network cluster is a collection of interconnected computers which share resources such as data storage. The individual computers, or nodes, are connected through both a physical and a software-level interconnect. The independent nodes are integrated into a single virtual computer, appearing to an end user as a single computing resource. If one node fails, the remaining nodes will handle the load previously handled by the failed node. This multiple computer environment provides many benefits to a user including high availability and increased speed of operation. (column 1, lines 30-40)

This paragraph describes a network cluster, but says nothing about a coordinator virtual device. The network cluster of the paragraph simply describes independent nodes of a cluster are integrated into a single virtual computer and appear to an end user as a single computing resource. The appellants respectfully submit that Frank's cluster as a whole neither teaches nor suggests the claimed "coordinator virtual device," which is not a cluster. The coordinator virtual device is used, for example, to assist in removing one of the nodes from a cluster when the cluster is partitioned and is clearly not the cluster itself. Frank further fails to teach or suggest an attempt to gain control of this network cluster as a whole.

Frank's shareable storage 22 is not a coordinator virtual device. The referenced portion of Frank states:

Here, shareable storage 22 has been illustrated as a single storage disk or the like. It should be understood by one of ordinary skill in the art that the shareable storage may include multiple storage devices. To implement multiple storage devices as the shareable storage 22, a header 25 of each storage device may include data indicating the identity of all devices comprising the shareable storage 22, a version number for information contained in the header 25, and any other pertinent data. To gain membership in the quorumless cluster 10, a node must have access to all storage devices comprising the shareable storage 22. (column 3, lines 34-45).

This portion of Frank merely describes shareable storage for use by cluster nodes. Although Frank does teach that to gain membership in the quorumless cluster, a node must have access to all storage devices comprising the shareable storage, Frank fails to teach or suggest that the shareable storage is a coordinator virtual device. Moreover, Frank's shareable storage is not something for which an attempt is made to gain control, as would be necessary to comport with the appellants' claims. In other words, mere access to a shareable storage device is not the same as control of the shareable storage device.

Frank's cluster manager 32 is not a coordinator virtual device. The Examiner states "Frank shows the use of a cluster manager, herein interpreted as a coordinator virtual device that oversees the removal and addition of nodes while considering whether or not the attempt to gain control of the coordinator virtual device is successful." Final Office Action of January 27, 2006, p. 12, ¶b. This argument is paradoxical in the sense

that the Examiner is suggesting the cluster manager oversees removal of nodes (which might include the cluster manager itself) while considering whether or not the attempt to gain control of *itself* (i.e., that which the Examiner equates with the coordinator virtual device) is successful. Moreover, this line of argument is also inconsistent with the teachings of Frank. Frank describes the cluster manager as follows:

Specifically, the cluster manager 32 manages cluster connectivity in the computer network cluster 10. For example, the cluster manager 32 can oversee the addition of nodes to and removal of nodes from the computer network cluster 10. It can also prevent the cluster 10 from partitioning into multiple cluster partitions. In addition, as an entity, such as an application or distributed lock manager, begins operating on node_1 12 and within the quorumless cluster 10, the entity may register with the cluster manager 32. Registration with the cluster manager 32 signifies that the entity requests that changes in cluster membership, among other things, be communicated to the entity by the cluster manager 32. (column 4, lines 31-42)

Note that cluster managers reside on each node. Column 4, lines 16-18. Frank fails to teach or suggest that the cluster manager is a virtual device, or a coordinator virtual device within the meaning of the term as used by the appellants. It is also unclear what way, if any, a cluster manager corresponds to at least a portion of a physical data storage device, as required by appellants claims. There is also nothing in Frank teaching or suggesting attempting to gain control of a cluster manager, either by a cluster manager itself, or any other entity.

Finally, Frank's coordinator node is not a coordinator virtual device. In the Advisory Action of April 24, 2006, the Examiner states: "Frank discloses if a coordinator node fails to operate during an update to the cluster definition, the remaining nodes of the cluster select a new coordinator node upon loss of the coordinator node" The referenced portion (column 8, lines 43-61) of Frank states:

The log file 54 and valid bits 51a-51d provide a means by which the quorumless cluster can recover if, during an update to the cluster definition 48, the coordinator node fails to operate. Upon loss of the coordinator node, the remaining nodes, node_1 12, node_2 14 and node_4 18, of the quorumless cluster 10 select a new coordinator node.

For illustrative purposes, it will be assumed that the remaining nodes designate node_4 18 as the coordinator node. Upon designation as coordinator node, node_4 18 checks the state of the repository 46 to

determine whether an update to the cluster definition 48 was left incomplete by the failure of the former coordinator node, node_3 16.

If a valid bit 51a-51d is set in the scratch area 50, the new coordinator node will examine the log file to determine whether an update had been started by the former coordinator node. If it was, the coordinator node parses the log file to determine where during the update process the former coordinator node failed. The new coordinator node completes the update from the identified point on.

Frank teaches that a coordinator node can update the cluster definition and apply the suggested changes. However, the coordinator node is still one of the nodes in the cluster. It is not a coordinator virtual device. Moreover, there is no teaching or suggestion of an attempt to gain control of the coordinator node. An existing coordinator node may leave the cluster, and a new coordinator node may be selected, but neither of these operations teaches or suggests attempting to gain control of either a lost or a new coordinator node.

Additionally regarding independent claim 15, Frank fails to teach or suggest a system including:

- a virtual device configuration server coupled to the first storage device and including a first memory and a first processor configured to provide a coordinator virtual device corresponding to at least a portion of the first data storage device;

- a plurality of virtual device configuration clients configured as a computer system cluster, at least one of the plurality of virtual device configuration clients including a second memory and a second processor configured to:

 - detect when the computer system cluster is partitioned;

 - attempt to gain control of the coordinator virtual device corresponding to at least a portion of the first data storage device; and

 - remove the at least one of the plurality of virtual device configuration clients from the computer system cluster when the attempt to gain control of the coordinator virtual device is unsuccessful.

For the same reasons as stated above, Frank fails to teach or suggest the claimed coordinator virtual device and the functionality to attempt to gain control of the coordinator virtual device.

Regarding the virtual device configuration server, the Examiner refers to Figure 5, and column 7, line 54 through column 8, line 15 which states:

Upon entering the proposed changes to the cluster definition 48 in the scratch area 50a, node_1 12 sets the valid bit 51a to indicate that it has completed entering its changes and notifies the coordinator node, node_3 16, that it has proposed changes to the cluster definition 48 at step 62. The coordinator node, node_3 16, verifies that the valid bit has been set and sets the update in progress flag 56 at step 64 to indicate that an update of the cluster definition 48 is in progress.

The coordinator node reads the scratch area 50a for the proposed changes at step 66 and increments the version number 58 of the repository to indicate a change is being made to the cluster definition 48 at step 68. At step 70, the coordinator node updates the cluster definition 48, to reflect the proposed changes. In addition, the coordinator node logs a progression of the update procedure in the log file 54. At step 72, the coordinator node clears the valid bit 51a and the update in progress flag 56.

As described above, the cluster definition 48 may be comprised of multiple copies. See FIG. 4. Updates to multiple copies may be accomplished in the manner described above where step 68 is implemented in parallel across all copies of the cluster definition at once.

In an alternate embodiment including multiple copies of the cluster definition, the coordinator node reads the scratch area 50a and updates a backup definition 48b (as identified by the map file 52).

This portion of Frank discusses coordinator node operation, but it is unclear from the Examiner's reference what element of Frank is believed to correspond to the appellants' virtual device configuration server. If it is the coordinator node itself, the appellants' respectfully disagree. Nothing in the cited portion of Frank teaches or suggests that a coordinator node is "configured to provide a coordinator virtual device corresponding to at least a portion of the first data storage device." The appellants are left to conclude that it is the Examiner's position that the node is *both* the claimed virtual coordinator device and the virtual device configuration server. Such an argument is inconsistent with the claimed separate limitations and is simply not supported by Frank.

Accordingly, the appellants respectfully submit that independent claims 1, 15, and 28 are allowable over Frank. Dependent claims 2-14, 16-27, and 29-31 depend from independent claims 1, 15, and 28, respectively, and are allowable for at least this reason.

CONCLUSION

The appellants respectfully submit that claims 1-31 are allowable for at least the reasons stated above. The appellants respectfully request that the Board reverse the rejections of these claims.

Respectfully submitted,

A handwritten signature in black ink, appearing to read 'E. Stephenson', with a long horizontal flourish extending to the right.

Eric A Stephenson
Attorney for Appellant
Reg. No. 38,321
512-439-5084
512-439-5099 (fax)

CLAIMS APPENDIX

1 1. A method comprising:
2 providing a coordinator virtual device corresponding to at least a portion of a
3 physical data storage device;
4 detecting when a computer system cluster, including a plurality of nodes, is
5 partitioned;
6 attempting to gain control of the coordinator virtual device; and
7 removing at least one of the plurality of nodes from the computer system cluster
8 when the attempting is unsuccessful.

1 2. The method of claim 1 wherein the providing a coordinator virtual device
2 corresponding to at least a portion of a physical data storage device further comprises:
3 selecting the at least a portion of a physical data storage device;
4 associating a physical description of the at least a portion of a physical data
5 storage device with a coordinator virtual device identifier; and
6 allowing at least one of the plurality of nodes of the computer cluster to access the
7 at least a portion of a physical data storage device via the coordinator
8 virtual device identifier.

1 3. The method of claim 1 wherein the providing a coordinator virtual device
2 corresponding to at least a portion of a physical data storage device is performed by at
3 least one virtual device configuration server.

1 4. The method of claim 3 wherein the at least one virtual device configuration
2 server is separate from the plurality of nodes of the computer cluster and wherein at least
3 one of the plurality of nodes of the computer cluster further comprises a virtual device
4 configuration client.

1 5. The method of claim 1 further comprising:
2 reading cluster membership information from the coordinator virtual device
3 corresponding to at least a portion of a physical data storage device.

1 6. The method of claim 1 wherein the detecting when a computer system cluster,
2 including a plurality of nodes, is partitioned further comprises:

3 reading, as performed by one of the plurality of nodes, cluster membership

4 information from the coordinator virtual device corresponding to at least a

5 portion of a physical data storage device; and

6 determining whether the cluster membership information indicates that the one of

7 the plurality of nodes is a current member of the computer system cluster.

1 7. The method of claim 1 further comprising:

2 writing cluster membership information to the coordinator virtual device

3 corresponding to at least a portion of a physical data storage device.

1 8. The method of claim 1 wherein the coordinator virtual device corresponding to
2 at least a portion of a physical data storage device further comprises cluster membership
3 information.

1 9. The method of claim 1 wherein the coordinator virtual device corresponding to
2 at least a portion of a physical data storage device is a coordinator volume.

1 10. The method of claim 1 wherein the detecting when a computer system
2 cluster, including a plurality of nodes, is partitioned further comprises:
3 monitoring a network coupled to each of the plurality of nodes for a heartbeat
4 signal; and
5 determining when the heartbeat signal is not present for a specified period of time.

1 11. The method of claim 1 further comprising:

2 retaining the at least one of the plurality of nodes in the computer system cluster

3 when the attempting is successful.

1 12. The method of claim 1 encoded in a computer readable medium as
2 instructions executable on a processor, the computer readable medium being one of an

3 electronic storage medium, a magnetic storage medium, an optical storage medium, and a
4 communications medium conveying signals encoding the instructions.

1 13. The method of claim 1 further comprising:
2 allowing at least one of the plurality of nodes of the computer cluster to
3 exclusively access the at least a portion of a physical data storage device.

1 14. The method of claim 1 further comprising:
2 obtaining exclusive access to the at least a portion of a physical data storage
3 device.

1 15. A system comprising:
2 a first data storage device;
3 a virtual device configuration server coupled to the first storage device and
4 including a first memory and a first processor configured to provide a
5 coordinator virtual device corresponding to at least a portion of the first
6 data storage device;
7 a plurality of virtual device configuration clients configured as a computer system
8 cluster, at least one of the plurality of virtual device configuration clients
9 including a second memory and a second processor configured to:
10 detect when the computer system cluster is partitioned;
11 attempt to gain control of the coordinator virtual device corresponding to
12 at least a portion of the first data storage device; and
13 remove the at least one of the plurality of virtual device configuration
14 clients from the computer system cluster when the attempt
15 to gain control of the coordinator virtual device is
16 unsuccessful.

1 16. The system of claim 15 wherein virtual device configuration server is further
2 configured to:
3 select the at least a portion of the first data storage device;

4 store a coordinator virtual device identifier associated with a physical description
5 of the at least a portion of the first data storage device; and
6 allow the at least one of the plurality of virtual device configuration clients to
7 access the at least a portion of the first data storage device via the
8 coordinator virtual device identifier.

1 17. The system of claim 15 wherein the first data storage device is at least one of
2 a disk drive, a JBOD, a disk array, and an integrated circuit.

1 18. The system of claim 15 wherein the first data storage device is coupled to the
2 virtual device configuration server via a network.

1 19. The system of claim 15 wherein the virtual device configuration server is a
2 volume server, wherein the coordinator virtual device is a coordinator volume, and the
3 plurality of virtual device configuration clients is a plurality of volume clients.

1 20. The system of claim 15 wherein the at least one of the plurality of virtual
2 device configuration clients is further configured to read cluster membership information
3 from the coordinator virtual device corresponding to at least a portion of the first data
4 storage device.

1 21. The system of claim 20 wherein the at least one of the plurality of virtual
2 device configuration clients is further configured to determine whether the cluster
3 membership information indicates that the at least one of the plurality of virtual device
4 configuration clients is a current member of the computer system cluster.

1 22. The system of claim 15 wherein the at least one of the plurality of virtual
2 device configuration clients is further configured to write cluster membership information
3 to the coordinator virtual device corresponding to at least a portion of the first data
4 storage device.

1 23. The system of claim 15 wherein the coordinator virtual device corresponding
2 to at least a portion of the first data storage device further comprises cluster membership
3 information.

1 24. The system of claim 15 wherein the at least one of the plurality of virtual
2 device configuration clients is further configured to retain the at least one of the plurality
3 of virtual device configuration clients in the computer system cluster when the attempt to
4 gain control of the coordinator virtual device is successful.

1 25. The system of claim 15 wherein the first memory and the virtual device
2 configuration server belong to at least one of a host computer system, a cluster node, a
3 storage appliance, a network appliance, and a storage area network (SAN) switch.

1 26. The system of claim 15 wherein the at least one of the plurality of virtual
2 device configuration clients is further configured to obtain exclusive access to the
3 coordinator virtual device.

1 27. The system of claim 15 wherein the virtual device configuration server is
2 further configured to allow exclusive access to the coordinator virtual device by the at
3 least one of the plurality of virtual device configuration clients.

1 28. An apparatus comprising:
2 a means for providing a coordinator virtual device corresponding to at least a
3 portion of a physical data storage device;
4 a means detecting when a computer system cluster, including a plurality of nodes,
5 is partitioned;
6 a means for attempting to gain control of the coordinator virtual device; and
7 a means for removing at least one of the plurality of nodes from the computer
8 system cluster when the attempting is unsuccessful.

1 29. The apparatus of claim 28 further comprising:
2 a means for reading cluster membership information from the coordinator virtual
3 device corresponding to at least a portion of a physical data storage
4 device.

1 30. The apparatus of claim 28 further comprising:
2 a means for writing cluster membership information to the coordinator virtual
3 device corresponding to at least a portion of a physical data storage
4 device.

1 31. The apparatus of claim 28 further comprising:
2 a means for determining whether cluster membership information stored in the
3 coordinator virtual device corresponding to at least a portion of a physical
4 data storage device indicates that the one of the plurality of nodes is a
5 current member of the computer system cluster.

EVIDENCE APPENDIX

None.

RELATED PROCEEDINGS APPENDIX

None.